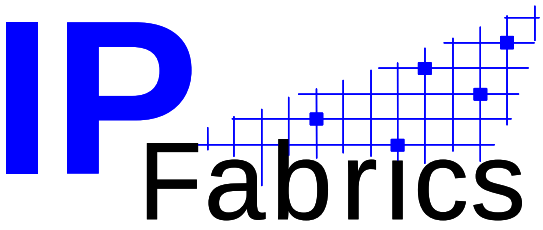

DeepSweep™

IAS Surveillance Modules

IAS Controller
IAS Content

User's Manual

October 2007

IP Fabrics Corporate Headquarters
14964 NW Greenbrier Parkway
Beaverton, OR 97006
Telephone (main line): 503 444-2400
Telephone (FAX line): 503 444-2401
Website: http://www.ipfabrics.com

Table of Contents

Table of Figures

Table of Tables

1 Introduction

This document describes the two Surveillance Modules for CALEA "broadband" use. These are referred to as:

- ias_controller_sm
- ias_content_sm

Detailed information on how to create SMs, SM actions, and creating surveillance assemblies (SAs) can be found in the DeepSweep™ User's manual. This document assumes those concepts are understood by the reader.

1.1 Implementation note

This document refers to several items that are not supported in the current release of the product. None of these restrictions are expected to limit the adherence of DeepSweep to the required standard but you should consider these aspects prior to initial deployment and work with IP Fabrics to ensure compliance.

- Support for Diameter protocol.
- Support for multiple communicating DeepSweep systems.
- Full support for IPv6 addresses.

1.2 Overview

This defines the architecture and external user interface of that part of DeepSweep that provides support for lawfully authorized electronic surveillance of Internet access and services. This support is consistent with the ATIS "T1.IAS" standard, now formally called ATIS-PP-1000013.2007.

DeepSweep provides broadband CALEA intercept completely self-contained. That is, it does not rely on separate CALEA support in routers, switches, access concentrators, RADIUS servers, etc., nor does it rely on separate probes and mediation/delivery systems.

The capability is provided in DeepSweep by two surveillance modules (SMs) as shown below. The rationale for providing the capability in two SMs is the following:

- The two SMs will likely listen on different networks in some ISPs
- Having multiple SMs allows the work to be spread over multiple PIXLs
- In some complex ISP networks, the two SMs may need to reside in different DeepSweep systems (future consideration)

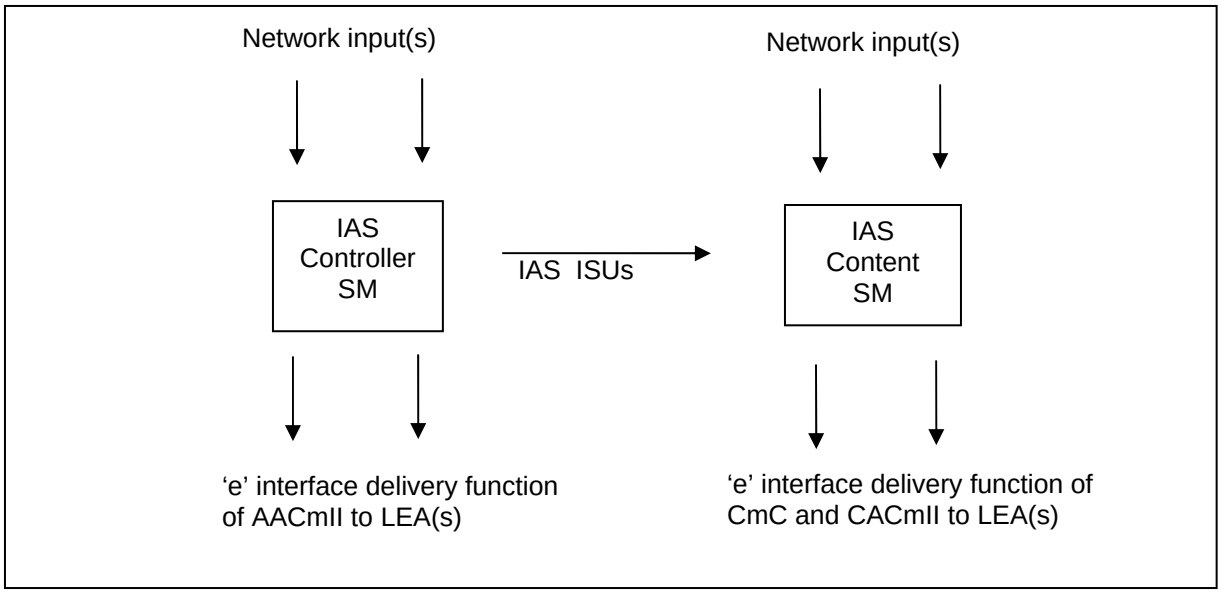

Figure 1. High level system showing relationship between SM types

In a nutshell, the controller SM watches registration and IP address assignment for the active intercept cases. For registration it understands RADIUS, Diameter (future), CHAP, and PPP Discovery. For IP address assignment it understands DHCP, IPCP, RADIUS, and Diameter (future). The content SM watches IP packet traffic for these cases. If content surveillance is authorized for a case, the content SM transmits the discovered packets as CmC. If content surveillance is not authorized, the content SM transmits CACmII messages.

Thus the chain that the controller SM is in must be connected to networks containing the registration traffic. The chain the content SM is in must be connected to networks where the subjects' packet data flows. There can be multiple content SMs per controller SM, and they can be in the same and/or different DeepSweep systems than the controller SM. (Use of multiple DeepSweep systems is a future capability.)

The architecture allows multiple intercepts to be active simultaneously; multiple intercepts on behalf of different LEAs (law enforcement agencies), including multiple intercepts on the same subject (aka subscriber or target); and adding or deleting intercept cases without interrupting ongoing intercepts.

1.3 Cases, Subjects, and Sessions

It is important first to understand the relationships among *cases, subjects, subject ids, access sessions,* and *packet data sessions*.

Case	Typically a court order authorizing surveillance, typically of a single subject.
Subject	A term used loosely herein. Typically a person to which a case applies. In T1.IAS the terms subject and subscriber are used interchangeably. There can be more than one case that involves the same subject.
Subject ID	A specific network identification of a subject. A subject can be known by multiple IDs, and thus a case can typically define multiple subject IDs. A subject ID can be associated with the person (e.g., logon name) or with equipment associated with the person (e.g., MAC address of laptop). Because subjects can be in multiple cases, so can subject IDs.

Access session	A state of a subject ID, where the subject ID has been registered or authorized for network access. Depending on the specifics of the network, a recognizable access session does not always exist, and when it does exist, it does not always have a recognizable termination point. Typically a subject ID has zero or one recognizable access sessions, although some networks permit multiple logons, in which case there are multiple access sessions.
Packet data session	A state of a subject ID where the user it represents can send IP traffic, meaning that an IP address has been assigned. Generally, if there is a recognizable access session, soon thereafter there is also one packet data session, but the converse is not true. There is at most one packet data session per subject ID when the subject ID is equipment, but there may be multiple packet data sessions (albeit not commonly) per subject ID when the ID is a name. A packet data session often does not have a recognizable termination point.

In DeepSweep, the two most important entities above are subject IDs and packet data sessions. Cases are more of an administrative concept, and access sessions may or may not exist.

2 Browser Pages

In DeepSweep all functions of the Surveillance Modules are configured via browser pages. The following describes the setup details for the pair of SMs for IAS.

2.1 IAS Controller configuration

Figure 2 shows the sole page for the controller SM. The left side shows the case information and the right side contains attributes about the controller SM as a whole.

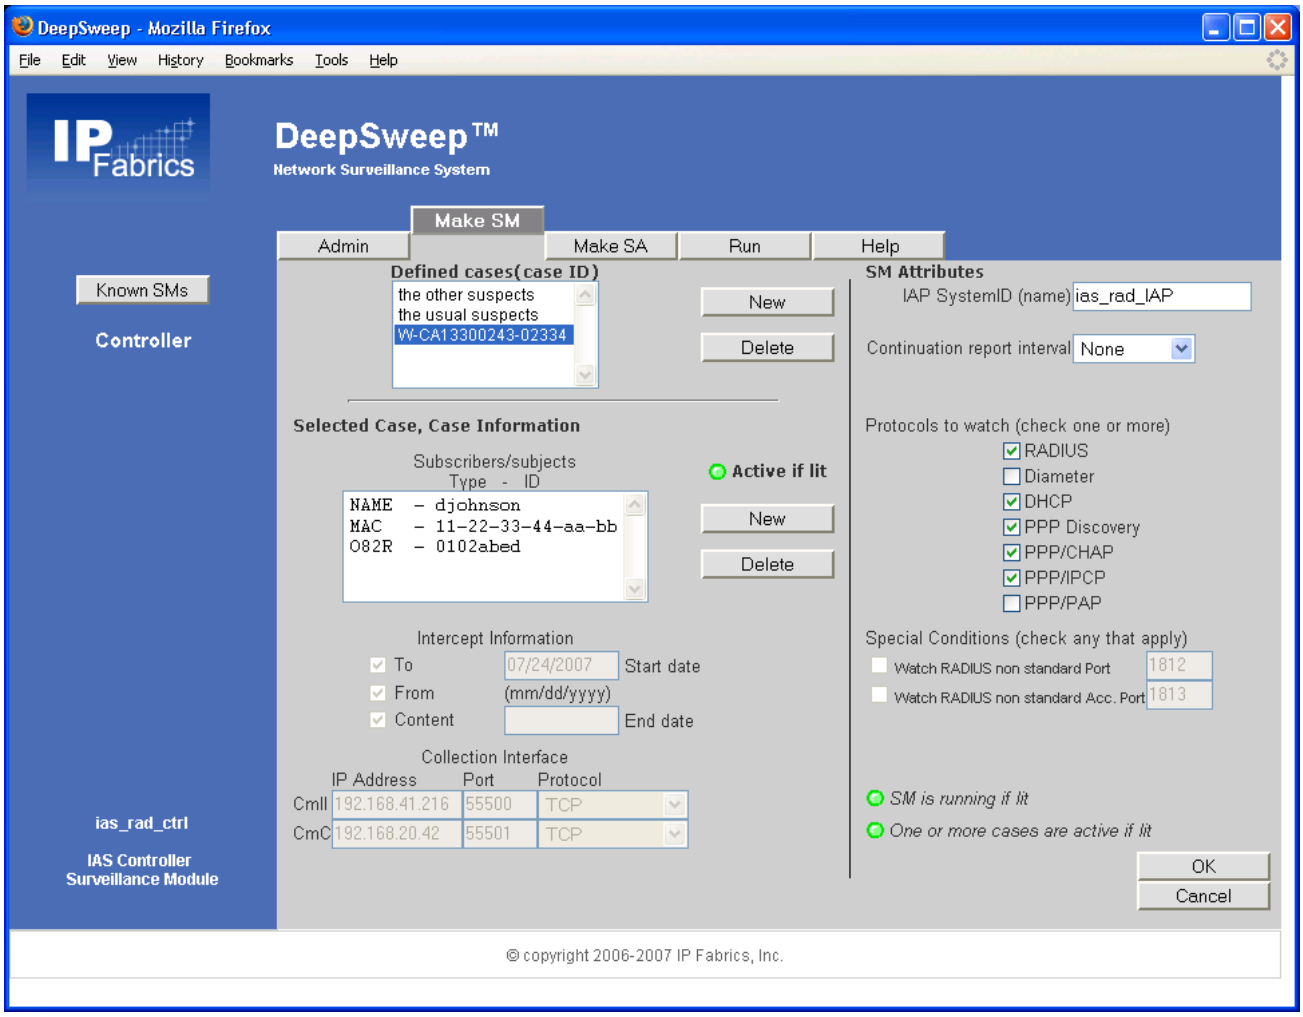

Figure 2. "IAS1" - Controller SM definition screen

The upper left has a scroll box that shows the name(s) of all cases that have been defined to this surveillance module. A case is identified by a case ID, which is a 1-25 character string. Depressing the NEW button brings up a box that asks you for a 1-25 character case ID. Providing that the case ID is different from all existing case ID's, depressing OK in that box brings you back to IAS1, where you can then describe the case beneath. Depressing DELETE next to the case scroll box causes the entire case to be deleted. If the case is not active, the system will prompt the user for confirmation. If the

case is active, the system will prompt the user with stronger wording, because deleting a case while it is active is unusual and serious.[1]

The bottom left side of the screen shows the definition of the selected case, or in the case of a newly created case, is blank. An indicator shows whether the case is active, meaning that surveillance is active. By definition, if this indicator is lit, the two indicators on the lower right are also lit.

A scroll box shows the subject IDs that are part of the case. We provide for any number of subject IDs per case because a subject may be known to the network in multiple ways. Subject ID can be a number of things, and the scroll box shows you the type that was chosen when the subject ID was entered. We provide the following types:

Type	Meaning	Notes and Examples
NAME	Name[2]	1-63 characters of any value. Can also be fully qualified domain name and network access identifier. We assume names are case insensitive. %HEXHEX can be used to insert UTF-8 or other encodings. For instance, the name x%C3%B3y puts the two hexadecimal byte values C3 and B3 between the characters x and y.
MAC	MAC address	12 hex digits. Hyphens may be used as visual separators. E.g., 00-5B-78-A4-00-9E
IPV4	IPv4 address (e.g., for static IP address assignment)	Must be standard dot notation e.g., 68.100.1.1
IPV4S	IPv4 subnet address	E.g., 62.100.0.0/16
IPV6	IPv6 address	We allow IPv6 colon notation, IPv6 compressed notation, and mixed notation (IPv4 address within an IPv6 address)
IPV6S	IPv6 subnet address	E.g., 805B:2D9D:DC28::/48
O82R	DHCP Option 82 remote ID	This can mean a lot of things and is very equipment specific. We interpret what is entered (see later entry screen) as hex digits, with hyphens allowed as visual separators
O82C	DHCP Option 82 circuit ID	This can mean a lot of things and is very equipment specific. We interpret what is entered (see later entry screen) as hex digits, with hyphens allowed as visual separators

Table 1. "SubjectID" type definitions

Depressing the NEW button next to the subjects scroll box brings up a box that asks you for the subject ID and its type. The type is selected from a set of choices in this box. Providing that the subject ID is different from all existing subject ID's of this type in this case and that its format matches the selected type, depressing OK in that box brings you back to IAS1. Adding a subject ID to an active case will cause that subject ID to be active immediately.

For all ID types, you fill in the identifier (e.g., name, IP address, MAC address) in the identifier box on page IAS2. Where the identifier is other than an IP address or subnet address, one can also optionally enter an IP address if a packet data session is already in progress by the subject at the start of the intercept. This IP address is treated as if it were a dynamically assigned IP address prior to the intercept.

[1] An active case is one for which surveillance is currently underway. To be active, the case must have at least one subject ID, must have its collection interface(s) defined, and the current date must be greater than the start date (if the start date is not blank) and not greater than the end date (if the end date is not blank).

[2] You will see later that we look for such names in multiple contexts – the obvious ones like RADIUS, but also places like DHCP option 61 client id and DHCP option 82 suboption 6 subscriber id.

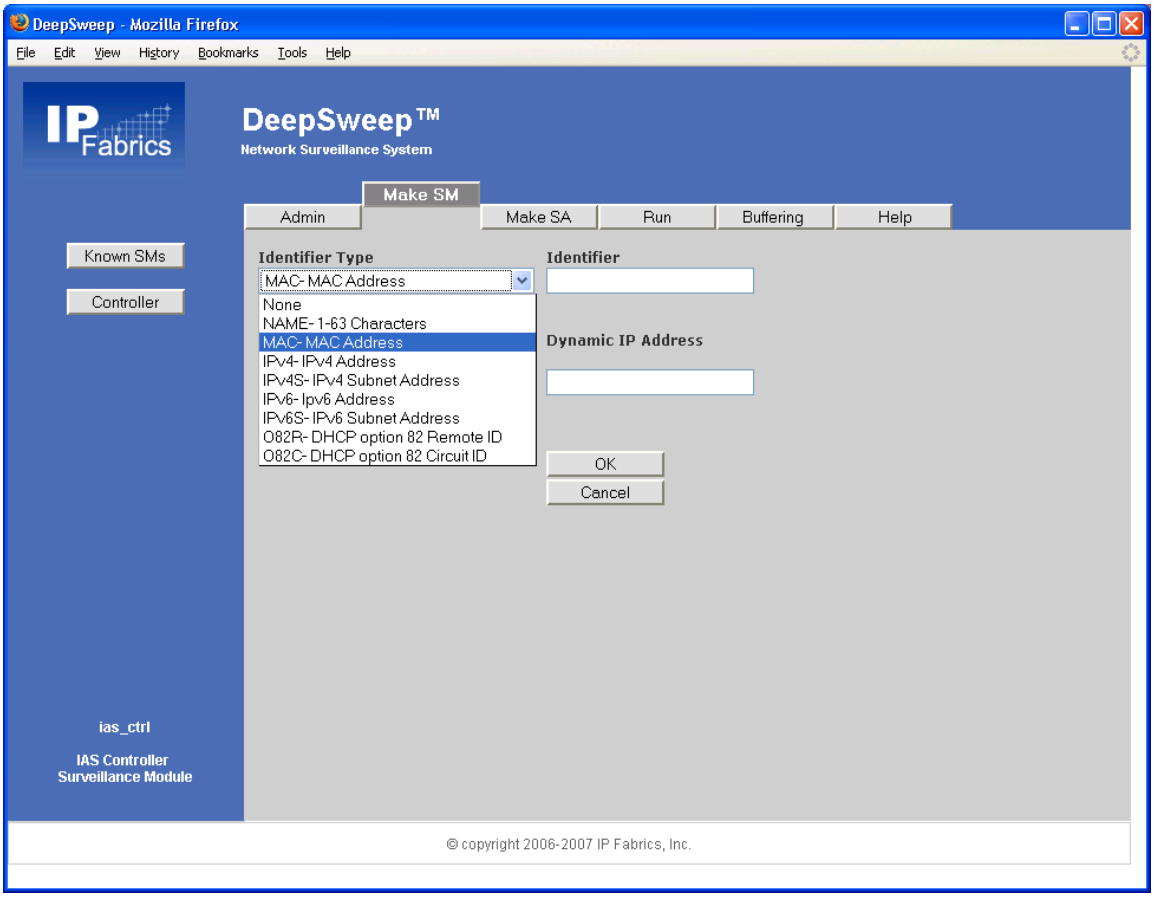

Figure 3. "IAS3" – New SubjectID definition

Returning to page IAS1, depressing DELETE next to the subject scroll box deletes the subject. If the case is active, the system will prompt the user for confirmation. Deleting a subject ID from an active case causes intercept related to that specific subject ID to stop for this case.

The next three check boxes define the type of intercept authorized.[3] At least one of TO and FROM needs to be checked. Note that even if content is not checked, one still needs to use the IAS content SM because it produces the CACmII messages about the packet traffic.

The next two fields are the dates, in the time zone of the DeepSweep system, on which intercept is to start and be completed. If start is left blank, it means "immediately." If end is left blank, there is no automatic cessation of the intercept.

The last piece of information is the collection interface to be used. The protocol and the CmII IP address and port are always required. If CONTENT is checked, the IP address and port of the CmC interface are also required. For the protocol, we provide four choices:

- UDP
- TCP
- UDP with appended message digest
- TCP with appended message digest

[3] T1.IAS doesn't distinguish between the "to" and "from" cases, but DeepSweep provides this capability – e.g., it distinguishes pen-register intercepts from trap-and-trace intercepts.

The last two address a section of the standard (5.3.5) that is peculiar because it states a "must" requirement but then only suggests a way to do it, and then never mentions the topic again in the remainder of the standard. The requirement is the ability to verify that the received CmII and CmC data was the same as the sent data. The suggestion is "a hashing method." But then this is never mentioned further. So, DeepSweep provides a way to provide it – if an appended message digest is opted for, we will compute a SHA-1 hash of the whole UDP or TCP payload and then append the first 96 bits of the 160-bit hash to the end of the payload, thus increasing the payload and packet size by 12 bytes.[4]

Now let's be more explicit about what it means to change an active case. The situations are the following:

- We delete a subject ID. We have discussed this earlier; intercept relative to this subject ID will stop.
- We add a subject ID. We have discussed this earlier also. The subject ID will become live immediately.
- We otherwise edit information on the page about the selected case. E.g., the to/from/content indicators, the collection addresses. The manner of attempting this would be to change the fields and click OK. This will be disallowed for an active case – either the fields will be unchangeable or the OK click will report an error.

Now to the information on the right of the page. First, at the bottom, note that this SM page is unusual in that it indicates if this SM is actually running as the page is viewed. Another indicator shows if any of the cases are currently active (intercept is active). The primary purpose of these are to help the user understand what adding or deleting an intercept will mean.

At the right are also some fields representing additional information, independent of specific intercepts, that is communicated when hits occur.

> IAP system id (name): A name denoting the intercept access point the DeepSweep system running this SM should be known by.
>
> Continuation report interval – selectable from 1 to 24 hours, or none. Continuation report interval is an optional message that we support from Annex C of T1.IAS.

Finally, there are a set of checkboxes for the user to give us a hint as to what protocols to listen in on. They are not really necessary in that we could just listen for them all, but if we know we don't need to listen for some types, we can run more efficiently.

If you select RADIUS as a protocol to watch then you can accept the standard ports (1812 and, for RADIUS accounting, 1813) or choose to enter non-standard ports if appropriate for your installation. For example, the older RADIUS ports may still be in use at your site.

2.2 IAS Content configuration

IAS2 (Figure 4) is the page for the content SM. We specify the name of the associated controller SM. We do it this way around so that there can be multiple content SMs associated with one controller SM, possibly on different DeepSweep systems, and multiple controller SMs, so we need to know the specific controller SM a content SM is related to. The IAP system id (name) has the same meaning as on the previous page. Whether content and controller SMs even in the same DeepSweep system have the same IAP id depends on the user and how the DeepSweep connects into his network.

[4] We use 96 bits because this is commonly done when producing an HMAC. What we define is not an HMAC (it is not computed with a secret key), but it would be easy to take the next step if necessary and make it a secure HMAC.

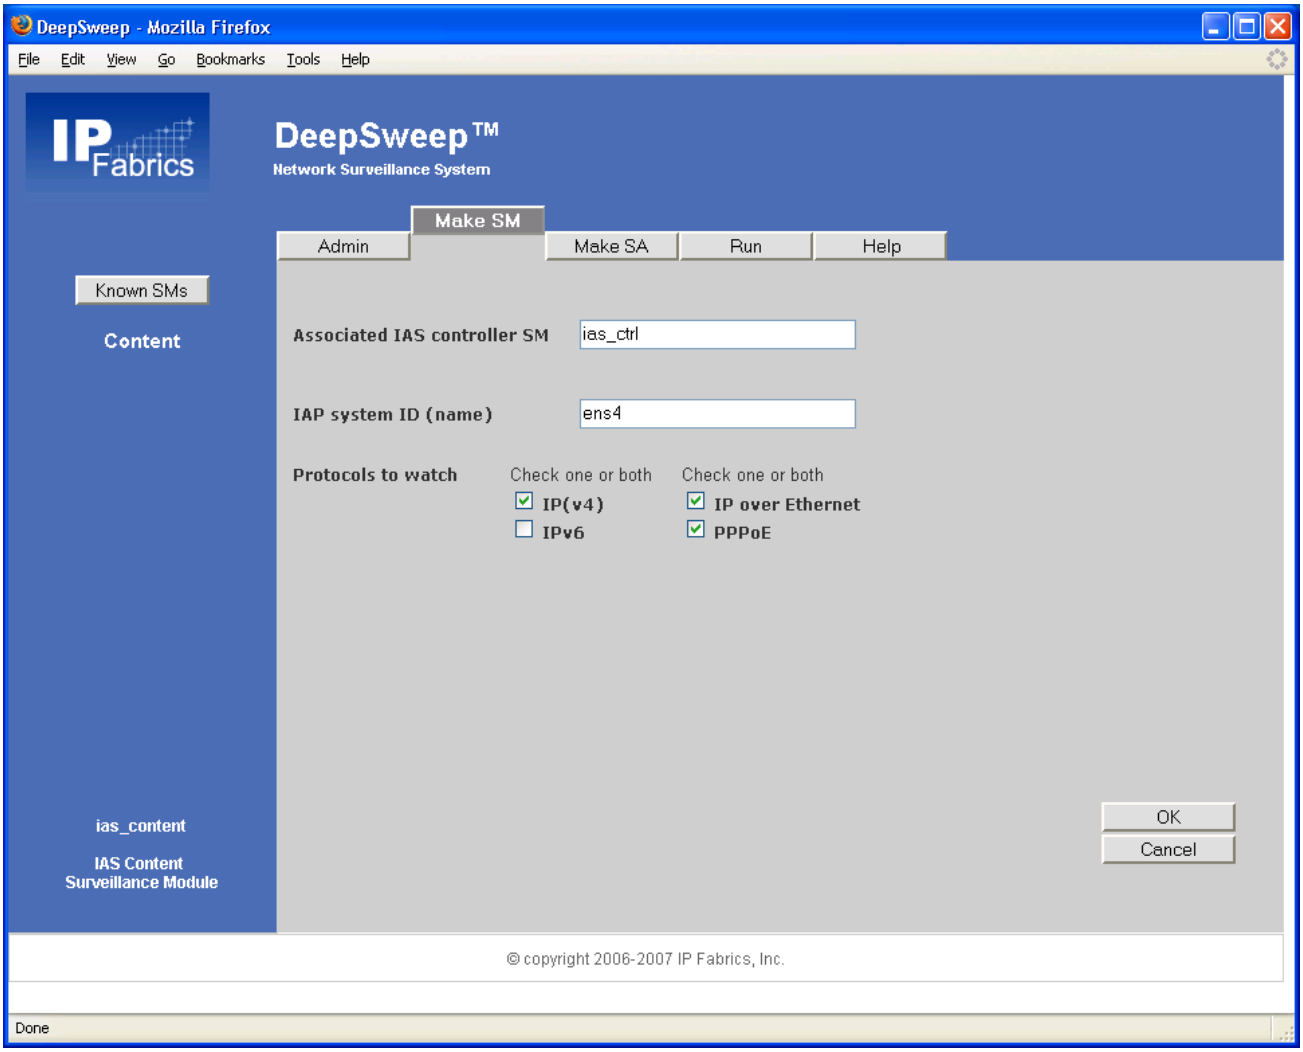

Figure 4. "IAS2" - Content definition screen

Regarding the packet traffic, one specifies whether to watch IPv4 and/or IPv6, and whether the IP is over Ethernet or over PPP over Ethernet.[5]

2.3 T1.IAS messages vs. normal DeepSweep "hit" actions

The IAS controller SM doesn't send the detected packet itself as an external message; it sends a T1.IAS CmCII message. Also it sends these messages in situations where there isn't a detected packet (e.g., at startup). The normal SM action options such as record, monitor, SMNP trap, kick to user program and reflect are not appropriate for lawful intercept. That said, Monitor and Record can be useful in confirming that an installation is working properly.

Up to this point, we've described the IAS SMs more in relation to what T1.IAS messages they generate rather than the usual DeepSweep description of "hits."

[5] What the user checks depends of course on the network topology and where the Ethernet for the chain in which this SM resides is tapped. Being able to exclude things we don't need to look at makes the SM run more efficiently.

It is unlikely that the IAS SMs will be used in chains with other SMs (other than perhaps using the controller and content SMs in the same chain in the case of a very simple network topology), but it is allowed. Unlike other SMs, the IAS SMs always pass both hits and misses on to the next SM by default when the chain is being defined and this is be the normal way to use the system.

Although it is valid to have in a system's surveillance assembly an IAS controller SM without an IAS content SM (in situations where one is using multiple DeepSweep systems), in most configurations it is an error. Hence, when one says run a surveillance assembly, if we see an IAS controller SM but no IAS content SM we will display a warning message but proceed.

2.4 SM Statistics

The RUN/STATISTICS page (described in the DeepSweep User' Manual) has a statistic for external messages sent. All CmII and CmC messages sent are counted in this system-wide statistic.

In the DeepSweep architecture, every SM can collect up to four statistical values that are represented on the RUN/STATISTICS page in a 2x2 matrix. The lower right corner is always the number of packets examined by the SM. For the T1.IAS controller SM, the statistics should be

Packets generating one or more AACmCII events.	AACmCII events
	Packets examined

For the T1.IAS content SM, the statistics are

Packets generating one or more CACmCII and CmC events.	CACmCII events
CmC events	Packets examined

3 Controller SM Logic

When the controller SM starts, it sends one surveillance-activation message per active case on the case's CmII interface. This message has the case id, the IAP system id, and the time.

There are two other circumstances during operation where a case becomes active. In both of these, a surveillance-activation message is sent:

- Time ticks from day N to day N+1 and the case's start date is day N+1
- A case is added and the current date is on or after its start date or its start date is blank, and the current date is on or before its end date or its end date is blank.

If the SA running is stopped, the controller SM sends one surveillance-deactivation message per active case. This message is also sent in the following two circumstances:

- Time ticks from day N to day N+1 and the case's end date is day N.
- An active case is deleted.

Periodically (the continuation report interval), the SM sends the surveillance-continuation message (unless the interval is nul). It sends one message for each active case.

All three surveillance- message types have the same information (case id, IAP system id, time). The surveillance-messages are considered CmII and are thus sent to the CmII interface of the associated case.

For all active subject IDs, the SM keeps some state information. A key piece of information is whether the subject ID has an IP address assignment. Some subjects can start with an IP address (the subject ID is the IP address). Some get one via a DHCP or IPCP assignment. Some can get one via a RADIUS or Diameter assignment or a notification to RADIUS in an accounting request. Note that IP address assignment is an important event; one key thing it triggers is notification of the content SM(s) of the IP address.

For a subject ID, its start date is the earliest start date of cases it is in, and its end date is the latest end date of cases it is in.

The controller SM watches RADIUS, Diameter (future feature), PAP, CHAP, and PPP Discovery requests and replies. Normally, a matching request or reply leads to the sending of one of five types of AACmII messages:

Access attempt
Access accepted
Access rejected
Access session end
Access failed

Note that here as well as in other situations, when the subject ID is in more than one case, we send a message per case.

3.1 RADIUS

If we discover the user name in a RADIUS access request, we will generate one or more access-attempt AACmII messages, containing the following:

Case id
IAP system id
Time

Subject id
Access method and access equipment id if specified
NAS id or IP address (comes from the access request)
Protocol signal = RADIUS

If we see a matching reply and its code is access accept, we will generate one or more access-accepted AACmII messages. This has the same information as in an access-attempt AACmII message, plus

- IP address (if the RADIUS server assigns an IP address).
- Access session identity. This is a unique value we create to denote a specific access session for the subject.

We also watch RADIUS accounting requests (code 4). These do not generate access-... messages, but could result in an IP address assignment. If we see a START accounting request, we see if the framed-IP-address attribute is present. If it isn't, we do nothing. If it is, if the subject ID is one we're watching, we take this as the IP address assigned to this subject ID by the NAS, and we do what is described later in the section on IP address assignment. If we see a corresponding RADIUS STOP accounting request, we assume this means unassignment of the IP address specified in the START.

If we see in a start accounting request the assignment of an IP address we're currently watching, we take this as an unassignment of the IP address.

3.2 Diameter

NOTE: Diameter protocol is not available in the current release.

3.3 PPP Discovery

PPPoE Discovery (Ethertype 0x8863) uses handshaking similar to DHCP but the goal is establishing a client / access concentrator link, not IP address assignment. However, in the case of a subject ID being a MAC address, PPP Discovery is watched for detection of an access session.

Since PPP Discovery is stateless, the sole things that need to be watched for, independently, are PADR and PADS packets. If we see a PADR whose source Ethernet address is a subject ID's MAC address, we have the access-attempt event. If we see a PADS whose Ethernet destination matches a subject ID's MAC address and if the PADS packet does not have an error TAG, we have an access-accepted event and we use the session ID in the PADS packet as the access-session ID. If the PADS has any error TAG, that denotes an access-failed event.

We will watch for a third thing as well. If we see a PADT packet sent to or from a subject ID's MAC address, we have an access-session-end event and, if there is an IP address, a packet-data-session-end event.

3.4 PPP CHAP/PAP

CHAP is an optional password authentication mechanism for PPP. CHAP is PPP protocol 0xC223. If we see CHAP, we will examine the challenge-response packet. In addition to the hash of the shared secret, the challenge-response packet contains the responder's name, which we will compare to the subject ID names. If we find a match, we have an access-attempt event and we will generate a new access-session ID. If we see a subsequent packet with the same CHAP identifier sent to the challenge-response sender, we see if it is a success or failure packet, which leads to the access-accepted or access-failed events. If we have an access-accepted event, we remember the MAC address for subsequent use with IPCP. PAP is also supported.

3.5 IP Address Assignment

Perhaps the most critical thing the controller SM does is track the assignment of IP addresses (or subnet addresses). This can happen in one of six ways:

- Via DHCP
- Via RADIUS
- Via NAS notifying RADIUS
- Via Diameter (not in initial release)
- Via IPCP
- Via the subject ID being defined directly as an IP address (or a subnet)

When an IP (or subnet) address is assigned, the controller SM needs to:

1. Remember this as state associated with the subject
2. Notify the content SM
3. Send one or more AACmII packet-data-session start messages

The message contains much of the standard stuff of the other messages, plus:

- Packet data session id. This is just a unique value identifiying the data session. Everytime we need a new one we increment a global counter.
- "IPAddress" (case is important). IPAddress is a data structure containing the IP address, the allocation method (e.g., static, dynamic), and optionally the IPv6 flow label and the prefix length.
- Some other optional items.
- If content intercept is authorized, a set of things: a unique content identifier so that content can be correlated to CmII, the address to where the content CmC will be sent, and an indication of which of the three alternate content forms will be used.

There is an AACmII message packet-data-session-failed when an IP address could not be assigned.

There is also an AACmII message packet-data-session-end. This gets sent if the subject's session is known to have been terminated (e.g., by Diameter). We generate this message in the following circumstance. If we believe an IP address is currently associated with a subject ID and we see a new assignment of the same IP address to anything other than this subject ID, we end the packet data session and send this message.

There is one other message associated with IP address assignment – packet-data-session-already-established. The intent is to cover the situation where surveillance is authorized for a subject who is already on the network communicating away. There is a footnote in the standard discussing "a number of approaches." The approaches cited are usually manual (e.g., look in the log on the RADIUS system to see if the subject is active). There is one situation where we send this message – when the controller SM is starting up and we see that a case/subject identifies the subject by an IP address or IP subnet address.

Both the packet-data-session-start and packet-data-session-already-established messages contain a delivery information parameter containing a set of things, some of them containing alternative choices. The system uses a unique content identifier, supplies the address of the content interface, and as discussed later, uses the ASN.1 delivery format as referenced in section D.2. in ATIS-1000013.2007. The "deliveryinformation" parameter is included only if content intercept is enabled.

3.6 Packet Data Session and Content IDs

We will use the same value for the packet data session ID and content ID. The value will be the four characters "*XXX*," where X is a numeral. In other words we will use the sequence 000,; 001,; 002,; ... Or as an octet string value the first one of the sequence is 3030302C.

3.7 DHCP

The IAS controller SM also examines DHCP requests and replies (UDP, ports 67 and 68). For this, it watches for any of the following conditions:

- Subject is MAC address and this matches ("hexadecimally") the chaddr parameter
- Subject is MAC address and this matches an option 61 client identifier
- Subject is name and this matches an option 61 client identifier
- Subject is name and this matches an option 82 suboption 6 (Cisco specific) subscriber id
- Subject is option 82 remote id and this matches ""hexadecimally") an option 82 suboption 2 remote id

The main message we want to watch is the DHCPACK, because this is the actual IP address assignment. However, because the option 61 and 82 parameters don't appear back in the DHCPACK, it is insufficient to look for just DHCPACK. We don't have to go back as far as DHCPDISCOVER and DHCPOFFER, but we do have to watch DHCPREQUESTs for the above cases. If we get a match, we remember the chaddr and the xid and if we see a DHCPACK with the same chaddr and xid subsequently, the IP address assignment in it is the one we want.

If we see instead a matching DHCPNAK or DHCPDECLINE, we generate the packet-data-session-failed AACmII message(s). We place the operation code in the "reasonForTermination" parameter.

If we see a DHCPRELEASE associated with an active subject (any of the four matching conditions above), we generate a packet-data-session-end AACmII message(s) (and notify the content SM to deactivate for this subject). If we see a DHCPACK that is assigning an IP address we are tracking, and the assignment is to other that the subject ID we think has the IP address, we treat this as a packet-data-session-end event for our subject ID.

There is also the situation where we have an active subject with an IP address, and we see a DHCPREQUEST from him. This would denote an attempt to renew. This is another unique state we track. There are three possible outcomes:

- If subsequently we see a matching DHCPACK with the same IP address as before, we take no special action as a result (i.e., we go on as if these hadn't occurred).
- If we see a different IP address assignment in the DHCPACK (which is an unusual case), we treat this as if a release occurred and then an assignment. I.e., we send packet-data-session-end and notify the content SM, and then we do what we do for a new IP address assignment.
- If we see a DHCPNAK, we send packet-data-session-end followed by packet-data-session-failed.

3.8 IPCP

IPCP is another protocol used to assign IP addresses. It is PPP protocol 0x8021. We will look for IPCP packets sent to any MAC address that either matches a MAC-address subject ID or the MAC address we recorded when we found a CHAP challenge-response that matched a subject name.

- If the IPCP packet is Configure-ACK to an IP address configuration option, we record the IP address value in the packet
- If the IPCP packet is Configure-Reject to an IP address configuration option, we have a packet-data-session-failed event.
- If the IPCP packet is Terminate-ACK, we have a packet-data-session-end event.

3.9 IPv6

IPv6 changes many of the rules of IP address assignment, eliminating the need for both DHCP and NAT. It remains to be seen how service providers handle addresses, but it is perfectly valid for a service provider to provide subscribers with just the ISP's IPv6 prefix(es), and let subscribers assign their own IPv6 addresses using the IPv6 stateless autoconfiguration. This protocol has nothing in it that identifies the subscriber, so it is not at all clear how the IPv6 subscribers will be discovered.

4 Content SM Logic

This SM is a lot simpler. It has a list of IP addresses and/or subnets that it should watch for, along with an indication of whether to watch for "to" (destination), "from" (source), or both, and whether content surveillance is authorized.

When we get a hit, we go down one of two paths depending on whether content surveillance is indicated for this case/subject-id:[6] [7]

- We send CmC (content) over the CmC interface.
- We send CACmII over the CmII interface.

When sending content, T1.IAS (Annex D) allows three choices of formats. It is sufficient to support just one. DeepSweep supports the ASN.1 delivery format as referenced in section D.2.

So, if we are intercepting content for this case/subject, we create and send a packet for each intercepted packet. The packet has an outer IPv4 or IPv6 header, with the source address being the DeepSweep and the destination address that specified for the CmC interface. We send it via UDP or TCP as specified by the user (IAS2) using destination port also specified there. The payload of the UDP or TCP packet consists of

- A packet header called the "CCDeliveryHeader".
- The encapsulated intercepted packet (starting with its IP header)

Note that this could result in a high percentage of fragmented packets being sent to the collection computer.

If content intercept is not enabled for a particular case/subject, the content SM still intercepts all IP packets to/from the subject's IP address and it provides header information from the intercepted packets. T1.IAS provides two ways to do this. DeepSweep does the following:

- Each packet results in the sending of a CACmII packet data header report message. This is the standard stuff (case id) plus the content id plus the two pair of IP addresses and ports.

T1.IAS standard does not specify for what transport protocols the ports should be provided. We provide them for TCP and UDP.

Note that when content is not to be sent for a particular case/subject, the header info is sent over the CmII interface, not the CmC interface.

If we are watching IPoPPPoE, we do a similar thing, except we need to recognize the PPP layer and when we see it is PPP protocol 0x0021, we look inside for the IP packet.

[6] In addition to the subject being part of multiple cases, where we need to send per case, one case might have content intercept authorized and another not, meaning we need to take different paths per case.
[7] The standard doesn't clearly say that CACmII is not delivered if CmC is delivered for a specific case-id/subject-id, but we assume it isn't.

5 Time

5.1 Time format

T1.IAS, like T1.678 and J-STD-025B, uses a timestamp in all CmII messages, and the ASN.1 syntax defines it as a type called GeneralizedTime, and this is a very specific format that ties the date and time to GMT or an offset thereof. Time is defined to the millisecond.

Once we detect a surveillance event, we immediately calculate a time stamp to be used in the message. Delay from detection to time-stamp value should be less than 1 ms. The delay from the event detection to the physical transmission of the message from the DeepSweep system must be less than 1 second 95% of the time.

5.2 Absolute Time Accuracy

For some reason, the T1.IAS standard does not spell out requirements for absolute time accuracy. For instance, it requires that the time placed in a CmII message be 200 ms or less from when the CmII event was triggered, but this is a relative requirement; it does not say that the time needs to be accurate to an absolute time standard. We will use NTP to keep DeepSweep's concept of time accurate to the NTP time standard; DeepSweep's time should be less than 200 ms inaccurate from absolute time.

6 Other CmII and CmC Interface Considerations

The transport protocol is not specified in the standard, but UDP and TCP are suggested. To provide a way of supporting the data integrity mentioned in section 5.3.5 of the standard, we calculate the TCP and UDP checksum on both interfaces and always transmit the correct checksum.

The only thing sent by the DeepSweep system to the LEA collection interfaces are CmII and CmC messages for the case(s) defined on browser page IAS1 as being associated with that specific interface. Note that when TCP is selected, this means that the normal TCP protocol handshaking packets also pass over the interface.

6.1 CmCII Messages

The table below summarizes the CmCII messages that we generate. For all messages, the parameters labeled M and O in the T1.IAS standard are always provided (the only O parameter is the protocol signal), and the parameters labeled C are provided when known or available (see earlier sections).

CmCII Message	Which SM	When	Notes
Access attempt	Controller	RADIUS/Diameter access request PPP Discovery PADR CHAP Challenge-response	
Access accepted	Controller	RADIUS/Diameter access accept PPP Discovery PADS CHAP success	
Access failed	Controller	RADIUS/Diameter access rejected PPP Discovery PADS with error CHAP failure	The rejection reason is placed in the message
Access session end	Controller	Diameter abort session answer Diameter session termination request PPP Discovery PADT	Not clear if ever generated this for RADIUS
Access rejected			
Access signaling message report			Not used. This is an alternative to the above five messages.
Packet data session start	Controller	RADIUS/Diameter access accept assigned an IP address DHCP DHCPACK IPCP Configure-ACK Controller SM startup when subject is an IP address Case goes active because start data occurs and subject is IP addr New case and/or subject ID is created where subject is IP addr	See earlier discussion of a number of special parameters created for this message
Packet data session failed	Controller	DHCP DHCPNAK DHCP DHCPDECLINE IPCP Configure-Reject	
Packet data session end	Controller	DHCP DHCPRELEASE PPP Discovery PADT IPCP Terminate-ACK Assignment of an in-use address	
Packet data session already established		One sent per subject ID per case on SM startup when subject ID is an IP address or subnet address	

Packet data header report	Content	Each intercepted subject packet	Only used if content intercept not enabled
Packet data summary report			Not used currently – possible future choice
Surveillance activation	Controller	One sent per active case on SM startup One sent when a case becomes active because start date occurs	
Surveillance continuation	Controller	Sent every N hours per active case	N is set on page IAS2
Surveillance change			Not used
Surveillance deactivation	Controller	SA termination – one sent per active case One sent when a case goes inactive because date ticks beyond end date One sent when an active case is deleted	

All CmII messages contain the case ID and IAP system ID, both of which come from information provided on page IAS1. All CmII messages also contain a timestamp, which is produced in real time (and see more information in section 5 on "Time"). And all CmII messages either contain either the subject/subscriber ID or, in the case of the packet-data-header-report message, the content ID.

The access-attempt and access-accepted messages contain two conditional parameters ("provide when known") – access method and network access node identity. The information in these two parameters is not accessible in any of the access protocols we watch for, so we treat this information as not known. That is also the case for the access-session-characteristics and location-information parameters in the access-accepted message and packet-data-session-start messages.

7 RFCs Supported

768 User Datagram Protocol
793 Transmission Control Protocol
1332 The PPP Internet Protocol Control Protocol
1994 PPP Challenge Handshake Authentication Protocol (CHAP)
2131 Dynamic Host Configuration Protocol
2132 DHCP Options and BOOTP Vendor Extensions
2373 IP Version 6 Addressing Architecture
2486 The Network Access Identifier
2516 A Method for Transmitting PPP over Ethernet (PPPoE)
2865 Remote Authentication Dial In User Service (RADIUS)
2866 RADIUS Accounting
2960 Stream Control Transmission Protocol
3046 DHCP Relay Agent Information Option
3174 US Secure Hash Algorithm 1 (SHA1)
3339 Date and Time on the Internet: Timestamps
3588 Diameter Base Protocol